



Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 112 (2017) 2506–2518

Procedia
Computer Science

www.elsevier.com/locate/procedia

International Conference on Knowledge-Based and Intelligent Information and Engineering Systems, KES2017, 6-8 September 2017, Marseille, France

Free Will and Turing Test with Multiple Agents:

An Example of Chatbot Design

Yoshiteru Ishida * and Ryunosuke Chiba

*Department of Computer Science and Engineering,
Toyohashi University of Technology
Tempaku, Toyohashi 441-8580, Japan*

Abstract

Free will is an important component of consciousness. Modeling an artificial consciousness requires clarifying the significance and the definition of free will. This study proposes a definition of free will similar to the epsilon-delta definition of continuous space (e.g., real numbers). Selection capability and frame expanding potential (i.e., the ability to allow for the exploration of further options), which are significant functions of free will, are discussed in the problem-solving context. We also propose a Turing test with multiple agents in which the intelligence of humans and machines will be relatively scored based on chats in a mixed community of humans and machines. Agents (machines) fail in the multiple agents Turing test because they lack the ability to evaluate chats with another agent, as well as chats between two other agents.

© 2017 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of KES International

Keywords: Chatbots; free will; artificial consciousness; epsilon-delta definition; Turing test; matching; mutual recognition model

* Corresponding author. Tel.: +81-532-44-6895; fax: +81-532-44-6895.
E-mail address: ishida@cs.tut.ac.jp

1. Introduction

In 2016, artificial intelligence (AI) defeated human experts in the game of GO. AI trained with big data has demonstrated competitive performance (and outperformed in cases with well-defined problems) in pattern recognition, participated in social networks, and even demonstrated the ability to create artistic works, such as paintings and musical compositions. The next question naturally raised is if AI should have consciousness^{3,4,6,11,12,17} or even free will^{1,2,10,14,15} in the future, or, in fact, if they might already have these characteristics if we restrictively define free will.

Nomenclature

$R_i(t)$	credibility (normalized to a continuous value from 0 to 1) of an agent i at time t .
T_{ij}	evaluation by agent i of agent j : 1 when agent i evaluates agent j as credible; -1 otherwise.

Among the components that underlie consciousness, we assumed that self-awareness may be formalized as the singular point used when mapping to form a world model⁸. We also considered that self-awareness, in that sense, can be used as an operating system (OS) for the self-related problems of robots. However, we also noted that this OS can face the frame problem in interactions with the world environment, including communication with other robots. This note proposes that free will may be a possible diversification mechanism, which will not only provide a solution to avoid deadlock and periodic interactions between two agents but will also allow for expansion of the frame of the world model, allowing further options when needed.

Aiming at free will as a component of a managing system, we try to define free will mechanically, which raises the further question of whether there is an objective method for testing if other intelligent entities, such as human and even machines, have free will or not. In pursuit of an objective test of free will, we propose a relative test with multiple agents.

Section 2 defines free will, aiming at determining if it exists in machines. Section 2 also reviews the significance of free will in a problem-solving context, focusing on its capability of expanding the world model and selecting options. Focusing on selection capability, Section 3 proposes a design for chatbots with a matching automaton. Section 4 proposes a relative Turing test by extending the Turing test¹⁸ with multiple agents to include humans and machines. The chatbots designed in Section 3 are used as machine agents, and the mutual recognition model will be used to score the test. Section 5 discusses the results and implications of the test. Section 6 discusses a design challenge.

2. Significance and definition of free will

2.1. Significance of free will

The subjective feeling of free will comes from the confidence that we can behave differently from what can be expected mechanically or deterministically by the environment outside the self. Aiming at building AI with artificial consciousness as a managing system, we are concerned with objective free will that can be tested with inputs and outputs. Further, we need an incentive to include free will as a component of a managing system.

The theory that humans may have originally harnessed consciousness with free will to avoid deadlock, or repeated interactions within a brief period (deadlock is a repeated interaction with one period), among individuals has evolutionary merit.

One of the significant functions of free will is to allow for selection that is free from external and past causation (Fig. 1). Another significant function is to allow for expansion of the world model, so that the expanded world model can provide further options for selection (Fig. 2).

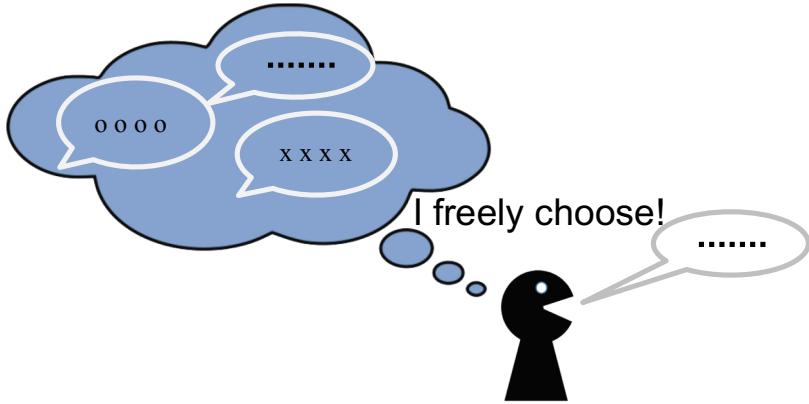
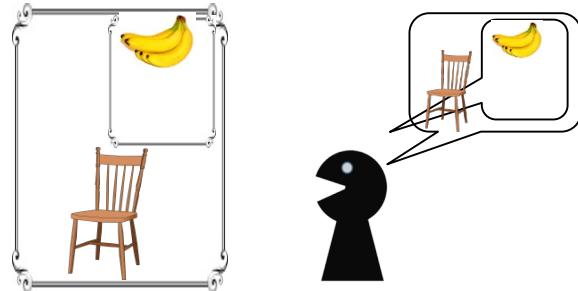


Fig. 1. Free will allows the agent to select an option, free from external and past causation.

Expanding the world model plays a significant role in problem-solving, since further options are made available. In Fig. 2, the agent wants to reach the banana. However, his only option is to grasp the banana, which is impossible because the agent is too short. By expanding the frame of the world model, he recognized another feasible option: grasping the banana while on the chair.



The idea of determining whether the target system has free will by observation is tempting. Let us create a thought experiment. Assume there is a system that takes input x and output $f(x)$ and that we do not know anything about how the target system generates $f(x)$. Aiming at testing whether the target system output $f(x)$ is based on some deterministic rule, we re-input the system output $f(x)$ again and again, and we observe the sequence of outputs: $x, f(x), f(f(x)), \dots, f(f(\dots f(x) \dots))$. If the target system is a closed system (without any other input channel) and a finite state machine, the output sequence will be periodic; however, we do not know anything about the target system. This thought experiment indicates the difficulty of defining and testing the target system merely by response observations. First, the generated sequence $x, f(x), f(f(x)), \dots, f(f(\dots f(x) \dots))$ could be aperiodic if the target system has free will, for if it recognizes x and $f(x)$ are the same in the first encounter, then it could choose otherwise in the next encounter after k iterations. The test of response periodicity has two merits: it can be easily implemented as a test and it can measure the degree of freedom to a certain extent. In fact, if $f(x)$ exhibits chaotic behavior, the period will be very long, and, hence, the chaotic functions are candidates that pass the aperiodic test. However, this test has

drawbacks; it can only indicate the possibility of free will because we know of functions that can generate an aperiodic sequence forever (e.g., computation of aperiodic infinite fractions, such as π and e), and testing for aperiodicity can require infinite sequences.

Free will, which is a significant aspect of consciousness, plays an important role in self-related problem solving⁸. Free will allows a problem-solving agent to choose options *freely* without any constraints. With the above thought experiment, the system with free will may choose freely among the multiple inputs x_1, \dots, x_k for the multi-input function $f(x_1, \dots, x_k)$, further the system will even expand input variables to $f(x_1, \dots, x_k, x_{k+1})$ by expanding the world model.

The significance of the problem-solving context is that symmetry breaking is possible when the problem is trapped in symmetric situations, which hamper problem solving. Humans, when trapped, can recognize the symmetric situation and will break the symmetry. We note that both symmetry recognition and symmetry breaking in self-related problem solving are significant acts of consciousness, namely self-awareness and free will.

2.2. Definition of free will

One can feel free will when one can react or can think in an unexpected manner. Thus, we believe current machines may not have free will because their interactions occur in a limited capacity that may be deterministically defined. Arguments on whether machines are capable of free will are similar to the arguments on whether we can define random numbers. Once we have defined free will, it would be difficult to admit that the machine behavior based on the definition indicates a free will. In this study, we try to define free will based on the tool used for the operational definition of the continuity of the real numbers, which is an epsilon-delta definition. Thus, free will may be formalized mathematically using an operational method.

The epsilon-delta method has been studied extensively for a long time, and it has been a basis for the fundamental study of infinitesimal analysis. Additionally, it has been the basis for the popular operation of differentiation and the definition of continuity of functions in mathematics.

Free will may be captured, at least from an observational point of view, only in an approximate manner. However, the approximated manner may include infinite repetitions, similar to the epsilon-delta definition of a limit in mathematics.

Let us define the system with free will as “any approximate system that can simulate the behavior of the target system, but can behave otherwise, against the simulated behavior,” where any interactions from outside are insulated. This definition reminds us of the epsilon-delta definition of continuity in mathematics.

With the above definition, the test to determine whether a target system has free will or not can be a non-deterministic problem for which only probabilistic statements are possible. It implies that finite state machines cannot simulate a system with free will. The definition also indicates that, without insulating the interactions (except stimuli), the system can generate an arbitrary behavior based on, for example, thermal noise.

If the system being tested is a finite state automaton that is interacting with a system that is also a finite state automaton, then the interaction between these two automata will ultimately converge on periodic interactions, even though the period can be long, depending on the number of states in these two automata. If the system being tested were to have free will (defined to have the capability of externally observing the system), then it would recognize the periodic interaction by observing from outside the system and avoid it by responding differently.

Suppose there is a machine that can correctly output the answer to the question, “What is the next prime number?” In the interactions between these two agents, there is no deadlock or periodic sequences. However, the machine (prime number generator) would not meet the definition of free will, for we can theoretically expect the output of the machine. However, if there is a random number generator that would output numbers different from those expected by any estimated distribution, the random number generator meets the definition of free will.

This property of free will-spontaneous symmetry breaking-could have promoted the evolution of free will in humans, since it has the evolutionary advantage of avoiding infinite fighting and breaking deadlock in a contest of survival of the fittest.

One possibility for a system that behaves as if it had a free will is an agent (machine) that is built as an open system with interactions from/to the internet. This note proposes a design based on a matching automaton⁷ that generates a decision based on the preferences among the constituent entities. We will use the example of chatbots

that operate based on sentences accumulated from a social networking service (SNS): Twitter. The matching automaton is a non-deterministic automaton that generates a decision (a stable matching of a set of pairings) based on the preferences of the two sets. Theoretically, it is identical to a solver of matching problems, such as the stable marriage problem.

3. Designing chatbots based on matching automata

3.1. Selection among options based on preference

This section discusses an exploration of agents (chatbots) that can interact (chat) with other agents naturally, as if they were human. The agent is built as an open system based on collections of chats from an SNS (*Twitter* in this case). Since we need many agents with different characters, we use the matching automaton to build many distinct characters.

Chatbots are designed based on the matching automaton, which requires two distinct sets (a set of agents and a set of sentences) and the preferences between the two sets (preference of each agent among a set of sentences and preference of each sentence among a set of agents). Each agent generates a sentence based on the preference, and each sentence is selected based on the preference. The algorithm for (the set of) agents to generate a sentence is as follows:

1. The data set of candidate sentences is prepared from the SNS.
2. Preference between the two sets (agent set and sentence set) is determined.
3. Stable matchings between the two sets are generated.
4. Agent optimal matching (as opposed to sentence optimal) is selected.
5. Agents post the selected sentence to the SNS.

Let us consider an example of chatbot design using a matching automaton.

Example 1. (Chatbots designed by matching automaton)

Let us consider an example of chatbot design by a matching automaton. In the chatbot design (Fig. 3), the preferences of the set of agents from among the set of sentences are determined based on several preference measures (e.g., length of sentence, number of formal forms (as opposed to casual forms), and number of specific word related to hobbies). When there are K preference measures, each agent a_i has K preference rankings for the set of sentences $\{t_i\}$ corresponding to the preference measures.

To characterize agents, L parameters are set for each agent. In this example, three parameters are set: p , patient (as opposed to short tempered); k , polite; and H : a set of words expressing the agent's interests, tastes, and hobbies. For example, if agent a_1 has the parameter, $p=3$, then a_1 's preference with respect to length of sentence is ordered based on how close the number of words in the sentence is to three. That is, patience p is reflected to the preferred sentence length.

The ranking from the set of sentences to the set of agents is determined similarly, which means the preference from sentences to agents is symmetric to the preference from agents to sentences.

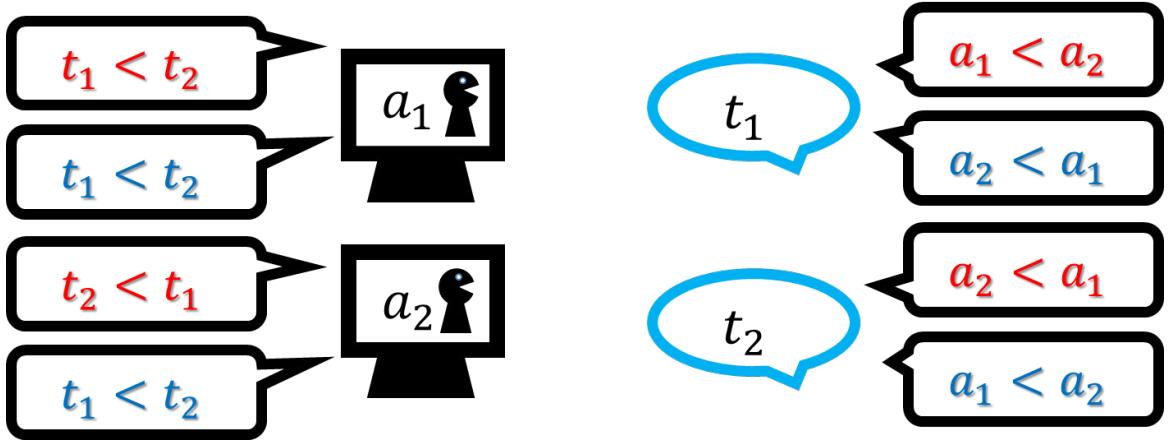


Fig. 3. Agents select a sentence based on stable matching between a set of agents and a set of sentences. Agent selects a sentence from the stable matching based on satisfaction.

Example 2. (Chatbot chats designed by matching automata and humans)

For this example, we designed and generated two different chatbots. Fig. 4 shows an example of chats among three agents: one human and the two chatbots, A and B .

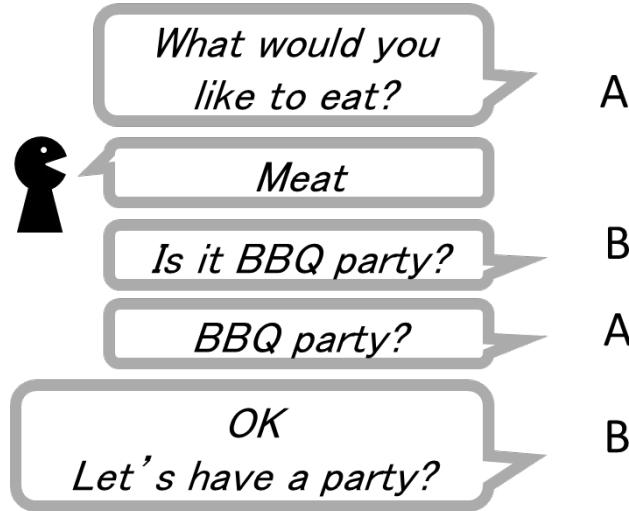


Fig. 4. An example of chats among three agents: one human and two chatbots, A and B.

3.2. Expansion of the world model

We can build agents that can chat with humans. However, humans with free will may not be satisfied with the given candidates for sentences. In these cases, humans may expand the world model by recalling past experiments. Currently, it would be difficult for agents (machines) to expand their world models because agents do not even know in which direction the world model should be expanded. What can be included easily when designing chatbots is the ability to search for another SNS or other internet media where more satisfactory sentences might exist. To achieve this, we (humans) need to prepare the mechanism for selecting the internet media (including a SNS). The chatbot must be able to choose the best available internet option (from its external memory, i.e., the internet).

4. Test for free will

Animals have two modes of decision making; one is automatic and reflexive, and the other is elaborative. If the decision-making situation requires quickness and there are few options, it will be reflexive. On the other hand, if the decision-making situation requires careful thinking and there are many options, it will be elaborative. Reflexive decisions often occur subconsciously, while elaborative decisions come up to conscious level.

To engineer free will in a problem-solving context⁸, at least two functions should be included: an expanding world model to create more options and the ability to select one of the available options. We have a feeling of free will when there are multiple options, and our decision is not totally dependent on past events, experiences, and the environment, and there is also the freedom to choose from other options. Although the capability of expanding the world model requires further study, one can test the agent's capability of selecting one option naturally, as if it were a person. This section focuses on this selection capability as a function of free will.

4.1. Multiple agents test with higher order recognition

The Turing test with multiple agents is not new⁵, even with respect to chat interactions¹⁶. The specific use of the Turing test proposed here involves the self/mutual recognition model⁹ (i.e., SRM or mutual recognition model), which is used to evaluate whether an agent is a human or a machine. This section introduces a SRM on which a dynamic network of agent credibility is constructed. The essential insight for using SRM for the Turing test with multiple agents is that an objective test should be used to perform relative tests between multiple agents.

The self-recognition model consists of nodes capable of recognizing the credibility of other nodes (i.e., credible, or not credible). The results of recognition are indicated by the arcs from recognizing nodes to being-recognized nodes and by the sign associated with the arcs (+ when recognized as normal and – when abnormal). Recognition by abnormal nodes is unreliable.

These self-recognition models can be mapped to a dynamic system called a dynamic relational network⁴ or a self-recognizing network¹⁰ with weighting and dynamic voting, where the weight and vote change dynamically through feedback from the changing vote. Weighting the votes and propagating them identifies the abnormal nodes correctly under certain conditions. A continuous dynamic network is constructed by associating the time derivative of the state variable (expressing the vote) with the state variables of other nodes connected by the evaluation chain. The vote is normalized to a continuous value (called credibility) ranging from 0 to 1 to show the inferred results as a generalization of the binary value (1 as true and 0 as false). Considering the effects from evaluating nodes and those from evaluated nodes, as well as retaining the intermediate information, leads to the following dynamic system, known as a gray model⁹:

$$\frac{dr_i(t)}{dt} = \sum_j T^+_{ji} R_j(t) - r_i(t)$$

$$R_i(t) = \frac{1}{1 + \exp(-r_i(t))}$$

where

R_i : credibility (normalized value of r_i).

r_i : credibility before normalization.

T^+_{ij} : $T_{ij} + T_{ji}$ if there is the arc from node i to node j or from node j to node i ; 0 otherwise (no arc).

T_{ij} : +1 (−1) for the arc from node i to node j with + (−) sign; 0 otherwise (no arc).

When evaluating nodes, node j will stimulate (inhibit) node i when $T_{ji} = 1$ (−1). We call this model the gray model, meaning that the network tries to determine the credibility of the node; namely, the credibility (which differs from the probabilistic concept of reliability) of a node becomes 1 (fully credible), 0 (not credible), or an intermediate value. Moreover, we propose different variants of this dynamic network, such as the skeptical model or the black and white model, for different engineering needs. The results of this note are generated only from the gray model because we need the detailed quantitative information that results from the weighting and dynamic voting (rather than binary results), as explained in Example 3 below.

Although credibility looks like the mathematical concept of probability, the only shared aspect is that the value is normalized from 0 to 1. Credibility does not have the mathematical rigor of probabilistic models, such as Bayesian networks⁸. For example, in mathematical models the probabilities of all exclusive events must add up to 1, while credibility does not consider the concept of exclusive events. For the computation of credibility, the only important point is consistency among the credibility of agents and evaluations between agents.

The self-recognition model can be used to evaluate whether each agent is human or not (machine). Capability of lying (machines imitating humans) could be a *singularity* that machines must achieve in order to seem human.

Example 3. (Objective Turing test of a machine with multiple agents)

As an example, let us consider a community of multiple (six) agents. The left side of Fig. 5 depicts the machine generated test results of the image on the right. As the signs attached to the arcs indicate, Agent 1 evaluates Agent 6 as a human, while Agent 6 evaluates Agent 1 as a machine. As a result of these mutual evaluations, Agent 2 earns the highest credibility, 0.634, as being human, and Agent 3 earns the lowest credibility, 0.010. In fact, this test was conducted with human agents (Agent 2 and Agent 5) and chatbots (Agent 1, Agent 3, Agent 4, and Agent 6). The mutual recognition network yielded correct answers. Agent 1, Agent 3, Agent 4, and Agent 6 failed the test. Agent 3 and Agent 4 had especially low credibility. The difficulty of this test is that chatbots also need to indicate whether the target agent is human or chatbot. In this example, chatbot simply evaluate those who use many words as human. This limited of evaluation capability of chatbot is too simple for the chatbot to imitate human. It should be modified, for example, so that chatbot evaluates human when the target agents can have common topic or interest.

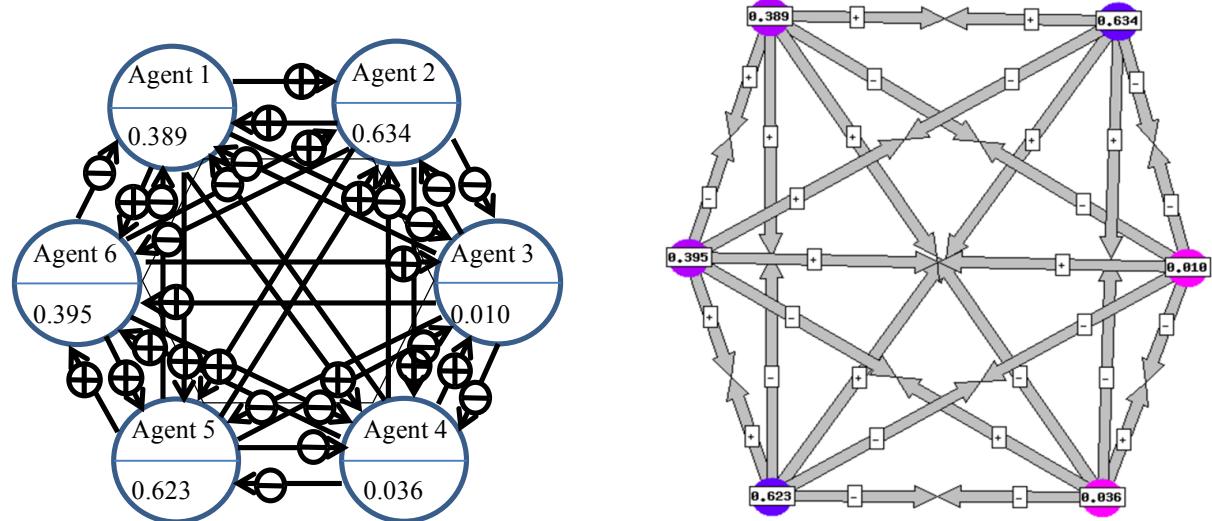


Fig. 5. Mutual recognition network for the Turing test with multiple (six) agents. Nodes correspond to agents, and arcs with the + (-) sign indicate if the source agent says the target agent is human (machine). Within the node, credibility of each agent is shown. Agent 2 (Agent 3) has the highest(lowest) credibility, and, hence, is evaluated as human (machine).

4.2. Experimental results of the multiple agent Turing test

In text communications on the Web, most dialogues are generated by multiple speakers. Therefore, we need to define the Turing test for multiple speakers. In the experiment (as shown in Fig. 6), the Turing test with multiple agents is performed as follows:

1. There are at least three agents.
2. Each agent must speak at least once.
3. Each agent judges whether the other agents are humans or machines.

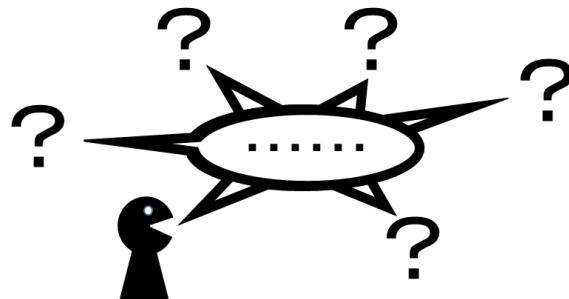
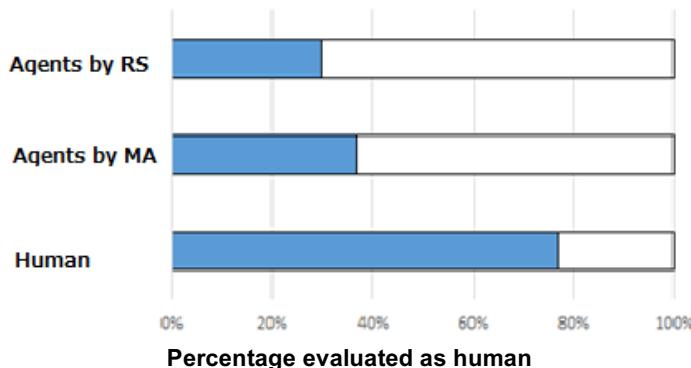


Fig. 6. Turing test with multiple agents. There should be at least three agents. Each agent must speak at least once. After chats, each agent judges whether other agents are human or machine.

Example 4. (Subjective Turing test by humans with multiple agents)

We conducted the Turing test with six agents, including at least one human agent. A total of 12 humans participated in the test. For a subjective Turing test by a human, we conducted a survey of the human agents who participated. In the questionnaire, we asked which of the other agents were human (not machine). Fig. 7 summarizes



the survey. Among the agents who participated in the test, humans were correctly evaluated as human with more than 70% accuracy. However, chatbots were identified as human at a rate of only about 30%. There is a slight difference between chatbots designed by a matching automaton (more than 30%) and chatbots that used randomly selected sentences (about 30%), but there is no statistically significant difference. However, there is a significant difference between the chatbots and the humans (about 40% difference). This indicates that the chatbots failed in the subjective Turing test, as well as objective Turing test (Example 3).

5. Discussion

Logically, we face a paradox because testing the intelligence of agents requires that agents possess the intelligence to test another agents' intelligence. However, for open systems that are influenced by the external environment, logical arguments may not be appropriate. In fact, we really do not know whether the intelligence of chatbots discussed in this study is artificial or human intelligence, for the sentences tested are from a SNS and communicated by humans.

As a byproduct of the multiple agent design (Section 3), it is known that as the number of agents increases, there appear to be higher modes of communication. That is, if only one agent is available, humans need a direct interaction with that agent. However, if two agents are available, human may not need a direct interaction with each agent (i.e., the human only needs to listen to the dialogue between the two agents). This kind of indirect and passive interaction may be required of a navigator agent in an automobile since humans need to pay attention when driving.

When conducting the Turing test with multiple agents (Section 4), the test turned out to be a test of whether agents can lie (i.e., pretend to be human). In fact, lying is an ability that even humans do not have until a certain age. Thus, agents capable of lying or pretending can be used for a game such as *Werewolf* (or *Mafia*)²⁰. In fact, the way that the wolf is evaluated is similar to the weighting vote used in the mutual recognition model.

6. A Challenge: How can machines expand the world model?

We defined free will to be consistent with the assumption that free will has a certain functional significance and can be built as a function of artificial intelligence with a certain mutual recognition capability. We avoided the philosophical and profound questions related to the definition of free will, which are beyond the scope of a single scientific field; these questions should be discussed while considering brain science, computer science, quantum physics, etc.

The essential point of this study is that we may reasonably assume that we can build agents that behave naturally enough so that average humans cannot discriminate humans from machines. With respect to the question of the definition of free will, we cannot know if it exists as a thing or phenomena before questioning if the thing (phenomena) is physical¹⁹ or biological⁶. Who can claim that free will could not be detected as gravitational waves are detected¹³, or that free will could not turn out to be just an illusion due to the human limited recognition of time (i.e., the feeling of (asymmetry in) time may be an illusion).

However, in the context of realizing one of the significant components of free will, the challenge is how to design a machine that is capable of expanding the world model, or even capable of introducing a new model into the conceptual map of the machine. That is a challenge that requires machines to also have inter-disciplinary conceptual maps.

7. Conclusion

In the pursuit of designing chatbots capable of interacting with humans so naturally that humans cannot discriminate between machines and humans, we focused on the significance of free will. We proposed that free will has evolved to avoid deadlock among agents (or between agents and the environment). As with the immune system, a spontaneous nature may be attributed to self-diversification, which can be used to avoid applying the same solutions to similar problems. The free will of artificial consciousness may be defined mathematically using the operational definition of epsilon-delta and may be tested relatively with multiple agents, including both humans and machines. We also proposed a Turing test with multiple agents, which allows testing to be done in a relative manner using the mutual recognition network.

As a component of artificial consciousness, the significance of free will is the ability choose among symmetrical options and to even enlarge the world model if there is no solution among these options. Consciousness as an operating system must contribute to problem-solving management, and artificial free will has an important role because the program used for problem solving could spontaneously exit the loop of command sequences. We conclude that, even for the restricted definition of avoiding deadlock in interactive behavior (between agents) or the

loop of problem-solving operations (between agents), the agent must be in an open system that requires some external noise for symmetry-breaking (similar to the immune system).

Acknowledgement

This note is inspired by WJ Freeman's pioneering approach to *Mind* using dynamical systems⁴.

References

1. Dennett D. Cognitive Wheels: The Frame Problem of AI. *The Philosophy of Artificial Intelligence* 1984; pp. 147–170.
2. Dennett D. What is Free Will? YouTube <https://www.youtube.com/watch?v=joCOWaaTj4A> (accessed on 2017/5/23)
3. Edelman G. *Wider Than the Sky: The Phenomenal Gift of Consciousness*. Yale University Press; 2005.
4. Freeman W. *How Brains Make Up Their Minds*. Weidenfeld & Nicolson: London; 1999.
5. Hingston P. A new design for a Turing test for bots. *IEEE Symposium on Computational Intelligence and Games* (CIG), IEEE, pp. 345–350, 2010.
6. Howkins J. *On Intelligence*. Henry Holt & Company: New York; 2004.
7. Ishida Y., & Hayashi T. Asymmetric phenomena of segregation and integration in biological systems: a matching automaton. In *Lecture Notes in Computer Science* 5712: pp. 789–796, 2009.
8. Ishida Y. A note on a continuous self-identification as self-awareness: an example of robot navigation. *Procedia Computer Science* 60: pp. 1865–1874, 2015.
9. Ishida Y., & Chiba R. Fermi estimate on the web: placing sensor networks on the Web with noise. *Procedia Computer Science* 60: pp. 1753–1762, 2015.
10. Kane R. (ed). *The Oxford Handbook of Free Will*. Oxford University Press; 2011.
11. Laureys S. (ed) *The Boundaries of Consciousness: Neurobiology and Neuropathology* 2006; **150**, Elsevier; 2006.
12. Koch C., & Tononi G. A Test for Consciousness. *Scientific American*; pp. 44–47, June 2011.
13. LIGO. LIGO Detected Gravitational Waves from Black Holes. <https://www.ligo.caltech.edu/detection> (accessed on 2017/5/23)
14. Ludowin D. Neuroscience & Free Will TEDx. YouTube <https://www.youtube.com/watch?v=DdoWRPzToag> (accessed on 2017/5/23)
15. Maye A., Hsieh C., Sugihara G., & Brembs B. Order in Spontaneous Behavior. *PLoS ONE* 2007; **2**:5 e443. 2007, <https://doi.org/10.1371/journal.pone.0000443>
16. Mauldin M. Chatterbots, tinymuds, and the Turing test: entering the Loebner prize competition. *Proceedings of AAAI '94*, 1994.
17. Tononi G. Consciousness as integrated information: a provisional manifesto. *Biological Bulletin* 215: pp. 216–242, December 2008.
18. Turing A. Computing machinery and intelligence. *Mind*, vol. 59, no. 236: pp. 433–460, 1950.
19. Wolf F. On the quantum mechanics of dreams and the emergence of self-awareness. In Hameroff SR et al. editors. *Toward a Science of Consciousness*. MIT Press; 1996. pp. 451–467.
20. Werewolf. <http://www.eblong.com/zarf/werewolf.html> (accessed on 2017/5/23)